

Measuring Patient Similarities via A Deep Architecture with Medical Concept Embedding

Zihao Zhu*, Changchang Yin*, Buyue Qian*, Yu Cheng[†], Jishang Wei[‡], Fei Wang[§]

*Xi'an Jiaotong University, Xi'an, Shangxi 710049, China

{zihaozhu, lentyery}@stu.xjtu.edu.cn

qianbuyue@xjtu.edu.cn

[†]IBM T. J. Watson Research, Yorktown Heights, NY 10598, USA

chengyu@us.ibm.com

[‡]HP Labs, 1501 Page Mill Rd, Palo Alto, CA 94304, USA

jishang.wei@hp.com

[§]Department of Healthcare Policy and Research, Weill Cornell Medical College, Cornell University

few2001@med.cornell.edu

Abstract—Evaluating the clinical similarities between pairwise patients is a fundamental problem in healthcare informatics. A proper patient similarity measure enables various downstream applications, such as cohort study and treatment comparative effectiveness research. One major carrier for conducting patient similarity research is the Electronic Health Records (EHRs), which are usually heterogeneous, longitudinal, and sparse. Though existing studies on learning patient similarity from EHRs have shown being useful in solving real clinical problems, their applicability is limited due to the lack of medical interpretations. Moreover, most previous methods assume a vector based representation for patients, which typically requires aggregation of medical events over a certain time period. As a consequence, the temporal information will be lost. In this paper, we propose a patient similarity evaluation framework based on temporal matching of longitudinal patient EHRs. Two efficient methods are presented, unsupervised and supervised, both of which preserve the temporal properties in EHRs. The supervised scheme takes a convolutional neural network architecture, and learns an optimal representation of patient clinical records with medical concept embedding. The empirical results on real-world clinical data demonstrate substantial improvement over the baselines.

Index Terms—Patient Similarity, Deep Matching, Medical Concept Embedding

I. INTRODUCTION

Patient similarity learning has been identified as one of the key techniques for healthcare transformation. During the past decade, Electronic Health events (EHRs), including diagnosis codes, lab results, prescription data, are becoming readily available for a huge amount of patients. This makes EHRs a valuable resource for evaluating the clinical similarities between pairwise patients. Patient similarity, which measures how similar a pair of patients are according to their historical information under a specific clinical context, will be the enabling technique for making various healthcare applications possible, such as cohort analysis, case based reasoning, treatment comparison, disease subtyping, and personalized medicine. In addition, learning patient similarity is a fundamental problem in evidence based medicine, which has been

identified as one of the major thrust areas for transforming healthcare and improving the quality of delivery of care.

Motivation. One of the key challenges to derive patient similarity measure is how to represent the medical events of patients effectively without loss of information. Since a great deal of healthcare analytics applications critically rely upon patient similarity, the similarity measures need to be both clinically effective and accurate. Though important, there are only a handful studies on patient similarity learning [1] [2]. Existing methods have successfully derived the similarity measure from EHRs data through mapping the medical events into vector spaces, however, their applicability is limited due to the lack of convincing explanations for patient representations in medical domain. There has been some existing work on applying patient similarity to various applications in medical literatures. However, there are still significant challenges on learning effective patient similarities, which, to our knowledge, have not been systematically addressed. (i) *Temporal-Sensitivity*: Temporal information is important to medical events, and is crucial to understand the dynamics of medical expressions. (ii) *High Dimensionality and Sparsity*: EHRs includes a wide range of data (such as diagnosis, medication, lab test) and a large number of possible medical events (over ten thousands of diseases and medications), so that EHRs data is usually represented in a high dimensional space. Besides, EHRs data is also very sparse, since a record exists if and only if the patient pays a visit to a specific clinical institute, for a particular condition. (iii) *Limited interpretability*: Due to the complexity of medical data, existing patient representation models are often weak at the perspective of clinical interpretations, which if addressed would significantly widen their applicability.

Proposal. Taking into account all challenges mentioned above, inspired by the idea of words embedding [3], we propose a method to represent patients and derive a similarity measure based on it. Unlike previous methods that model each medical event as a binary event vector over time (one if the medical event happened and zero otherwise), we derive

a fixed-length vector representation from EHRs by medical concept embedding. In text mining, a particular word can be predicted based on the context around it [4] [3]. Similarly, events happened before and after a specific medical event can be viewed as its medical context, which may be used to make event predictions in medical domain. Based on the medical context, each event is compressed into a given length vector with medical concept embedding. Similar to the word embedding [3], the event embedding presented in our model hold its natural medical concept. Furthermore, we adjust the range of context, with respect to the specific conditions of a medical event, to achieve an event embedding with temporal information. By stacking all event embedding vectors together, each patient is then represented as an embedding matrix. Note that, compared to describing patients using binary event vectors, the embedding extracts clinical features of a patient from EHRs and represent them in a reasonable dimension, resulting a natural dense embedding matrix for every patient.

Based on the embedding matrix representation of patients, we propose two methods, supervised and unsupervised, to derive the similarity measures. Note that the number of medical events varies from patients to patients, and both the supervised and unsupervised approaches are required to measure the similarity between matrices with different dimensions. As for the unsupervised method, we adopt the RV coefficient [5] and $dCov$ coefficient [6], respectively, to measure linear and non-linear relations between pairwise patients based on the embedding matrix. In the supervised model, we measure the patients similarity using the Convolutional Neural Network (CNN), where the deep medical embedding is obtained from the intermediate convolutional feature maps. With the given number of convolutional filters, an event embedding matrix is mapped to a fixed-length feature vector. The deep medical concept embedding contributes to improved patients similarity measures. We shall later in the paper make a comparison amongst different types of patient representations, including the binary event matrix representation.

Empirical Study. Patient cohort study is the most effective way to analyze the causes, treatments, and outcomes of diseases. To evaluate the representations we proposed, we conduct a cohort analysis based on the obtained measures of patient similarity. Our model is tested on a real-world EHRs dataset containing a wide range of medical events over a long time period. The experimental results demonstrate the effectiveness of our model in measuring patient similarity.

Contributions. Our work makes the following distinctive technical contributions:

- We adopt a state-of-the-art distributional representation model to project medical events to fixed-length vectors, which are then used to measure patient similarity.
- We effectively extract the low-dimensional and dense representation for patients from EHR data, with the temporal information preserved.
- We propose two solutions for patient similarity Learning, unsupervised and supervised. This makes our framework applicable to most similarity-related applications in

healthcare analytics.

The rest of this paper is organized as follows. Section II introduces related studies. Details about our model are presented in Section III. The experimental results are reported in Section IV, and Section V concludes.

II. RELATED WORK

In this section, we first review some related work on evaluating the clinical patient similarities, and then review some relevant problems associated with deep learning.

A. Patient Similarity

In healthcare informatics domain, there are a lot of works focusing on patient similarity. For example, [7] proposed a patient similarity algorithm named SimSvm that uses Support Vector Machine(SVM) to weight the similarity measures. [8] proposed a patient similarity based disease prognosis strategy named SimProX. This model uses a Local Spline Regression (LSR) based method to embed these patient events into an intrinsic space, then measure the patient similarity by the Euclidean distance in the embedded space. These methods do not take the temporal information into consideration when evaluating patient similarities. Wang [9] presented an One-Sided Convolutional Matrix Factorization for detection of temporal patterns. Cheng [1] proposed an adjustable temporal fusion scheme using CNN-extracted features. Based on patients similarity, plenty of applications are enabled. In [10], Ng provided personalized predictive healthcare model by matching clinical similar patients with a locally supervised metric learning measure. [11] proposed Integrated Method for Personalised Modelling (IMPM) to provide personalised treatment and personalised drug design.

There are many research have been conducted on clustering patients based on machine learning. In order to rate patients health perceptions, Sewitch [12] make cluster analysis using k-means to identify the patients groups based on the discovering the multivariate pattern. To capture underlying structure in the history of present illness section from patients EHRs, Henao [13] proposed a statistical model that groups patients based on text data in the initial history of present illness (HPI) and final diagnosis (DX) of a patients EHRs. For human disease gene expression, Huang [14] presented a new recursive K-means spectral clustering method (ReKS) to efficient cluster human diseases. Most of these research have demonstrate effectiveness of their model with real-world experiments, that convinces us of the applicability of clustering patients on cohorts discovering.

B. Embedding Learning and Semantic Matching

One of the most important components in our patients similarity measure is deep distributional medical concept embedding. Distributional representations has gone through the long evolution, and shows state-of-the-art results in many fields recently. [3], [4] proposed continuous Bag-of-Words model and Skip-gram model to represent words in vector

space. The word representations using neural networks provide state-of-the-art performance on measuring syntactic and semantic word similarities. Many works as well as ours are inspired by the words embedding with neural networks. [15] learned image embedding by concatenating skip-gram linguistic representation vectors with visual concept representation vectors. [16] encoded a query-document pairs into discriminate feature vectors using distributional sentence model. Our model achieves the goal of embedding patients clinical features in the dense matrices with modest dimensionality.

With medical concept embedding, we look forward to calculating the similarity amongst patients according to their EHRs. Considering the representations of patient medical events do not have a common time dimension, we cannot compare the patient event matrix directly. [17] provided a relevant similarity measures between temporal series of brain functional images belonging to different subjects. Similar to [17], we adopt the RV coefficient to measure patient similarities. Note, however, that this coefficient only considers linear relationships between two data sets. To do more systematic research on measuring similarity of patient, our model also measures non-linear correlation between two patients using dCov coefficient. Apart from those unsupervised approach, we adopt the supervised learning method. We modify the Convolutional Neural Network(CNN) to derive the similarity scores for pairs of patients. The Convolutional networks models which are originally invented for image processing have wide applications in other domains. [16], [18] and [19] respectively obtains the continuous representations of the sentences or short texts by a convolutional deep network, then the similarity can be effectively established.

III. THE PROPOSED METHOD

Accessing patients similarities in EHRs data is a very challengeable task. In this section, we will first propose to learn the contextual embedding of each medical concept. Then we provide an unsupervised method to estimate the similarity score, which takes the learned medical concept as the input. After that we exploit an architecture building with convolutional neural network to measure the similarity of pair patient records with some supervision encoded.

A. Contextual Embedding of Medical Concepts

Our goal in this step is to get the contextual embedding of each medical concept from patient EHRs, which provide a better representation for medical concepts than general one-hot encoding representation. By “context around a medical concept A we really mean the medical events happening before and after A within the patient EHRs corpus. For each patient, by concatenating all medical events in his/her EHRs according to their happening timestamps (for events with the same timestamp we do not care about the order), we obtained a paragraph describing the historical condition of him/her. So the context around a specific medical event is similar to the context around a word in a paragraph. How to derive effective word representations by incorporating contextual information is a

fundamental problem in Natural Language Processing (NLP) and has been extensively studied. One recent advance is the Word2Vec technique [4] that trains a two-layer neural network from a text corpus to map each word into a vector space encoding the word contextual correlations. The similarities (usually cosine distance) evaluated in such embedded vector space reflect the contextual associations (e.g., words A and B with high similarity suggests they tend to appear in the same context).

In NLP, the context around each word is usually identified as the adjacent words before and after it. In Word2Vec such context is defined by a sliding window around each word and the length of the window reflects the scope of the context. In EHRs, as there is a timestamp associated with each medical concept, we do not just want to consider the relative positions when defining the context, but consider the actual timestamps. For example, we may want to treat event B happened one year after A differently comparing to event B happened one week after A. Another factor we need to consider is the context scope around each event, i.e., the length of the sliding windows. In Word2Vec models for NLP every word is assigned with the same window length. In contrast, for EHRs, we may want medical concepts related to chronic conditions to have larger scopes while acute condition concepts to have smaller scopes. Moreover, because of the variabilities among individual patients, the scope for the same event could be different for different patients. Therefore we propose an adaptive way to determine the window length for an event in the EHRs of a specific patient. Our heuristic is that chronic conditions are more likely to appear repeatedly in a patients EHRs and thus have higher frequency, and acute conditions will be less frequent. Then for medical event i and patient p ,

$$L(i, p) = f(i, p) * a + \theta \quad (1)$$

where $f(i, p)$ is the frequency of event i in the EHRs of patient p . a and θ are constants.

B. Temporal Patient Representation

After the medical concepts embedding step, We expect that the medical concept representations learned by Skip-gram will show similar properties so that the concept vectors will support clinically meaningful vector additions. A straightforward representation of a patient p will be as simple as converting all medical concepts in his medical history to medical concept vectors, then summing all those vectors to obtain a single representation vector. However, this representation will lose the temporal information. Instead, we utilize a temporal representation: the records of each patient p is represented as a matrix \mathbf{X} with dimension $d \times N_p$, where d is the fix embedding dimension and N_p is the total number of visit patient p has. A single representation vector of one visit is obtained by summing all the medical vectors in that visit. Usually, N_p varies from patient to patient. Given two patients p_a and p_b , calculating the similarity between the record \mathbf{X}_a and \mathbf{X}_b is not that intuitive. We propose the method in the following sections.

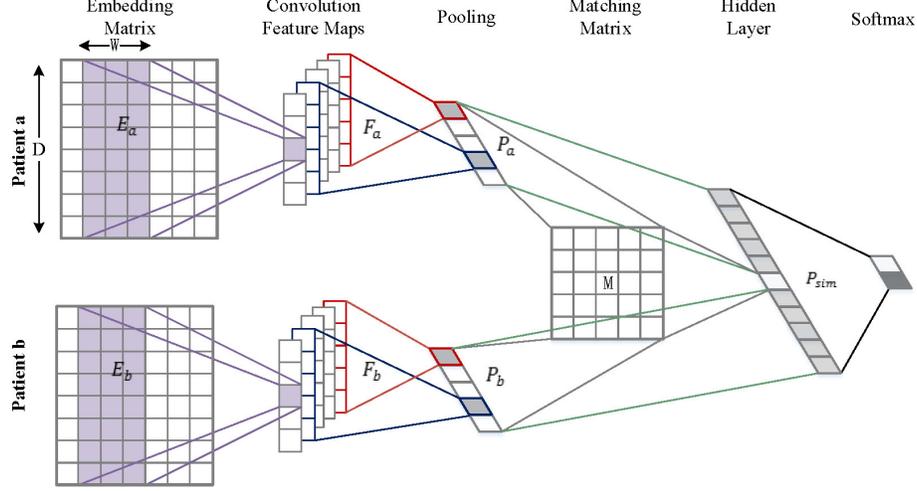


Figure 1: The overall framework of supervised patient similarity matching. To train the singular neural network, embedding matrices of pairs of patients E_a, E_b passed through convolutional filters are mapped into feature maps. We build the deep embedding patients representations P_a, P_b for patients by pooling patients feature maps into the intermediate vectors. With the rich feature vectors we learn a symmetrical similarity matrix M for measuring the distance between patient a and b .

C. Unsupervised Patient Similarity

In order to calculate the similarity score based on the patient temporal representation, we provide two alternatives. The first one is to utilize RV coefficient [] and dCov efficient to estimate the similarity over the pair of temporal patient representation. In particularly, given two matrix representations $\mathbf{X} \in \mathcal{R}^{n \times k}$ and $\mathbf{Y} \in \mathcal{R}^{m \times k}$, the RV coefficient is defined as:

$$\text{RV}(\mathbf{X}, \mathbf{Y}) = \frac{\text{tr}(\mathbf{X}\mathbf{X}'\mathbf{Y}\mathbf{Y}')}{\sqrt{\text{tr}(\mathbf{X}\mathbf{X}')^2\text{tr}(\mathbf{Y}\mathbf{Y}')^2}} \quad (2)$$

For the dCov efficient, let's first define the empirical distance covariance:

$$\text{dCov}_n^2(\mathbf{X}, \mathbf{Y}) = \frac{1}{n^2} \sum_{i,j=1}^n (d_{ij}^{\mathbf{X}} - d_{i.}^{\mathbf{X}} - d_{.j}^{\mathbf{X}} + d_{..}^{\mathbf{X}})(d_{ij}^{\mathbf{Y}} - d_{i.}^{\mathbf{Y}} - d_{.j}^{\mathbf{Y}} + d_{..}^{\mathbf{Y}}) \quad (3)$$

where $d_{ij}(\cdot)$ is the Euclidean distance between sample i and j of random vector \mathbf{x}_i , $d_{i.} = \frac{1}{n} \sum_{j=1}^n d_{ij}$, $d_{.j} = \frac{1}{m} \sum_{i=1}^n d_{ij}$, $d_{..} = \frac{1}{n} \sum_{i,j=1}^{m,n} d_{ij}$. The empirical distance correlation (dCor efficient) is defined:

$$\text{dCor}_n^2(\mathbf{X}, \mathbf{Y}) = \frac{\text{dCov}_n^2(\mathbf{X}, \mathbf{Y})}{\sqrt{\text{dCov}_n^2(\mathbf{X}, \mathbf{X})\text{dCov}_n^2(\mathbf{Y}, \mathbf{Y})}} \quad (4)$$

D. Measure Similarities with Supervision

In order to add some supervision to this procedure, we proposed a deep learning model. The idea is derived from semantic matching problem in NLP, which aims to determine a matching score for two given texts. Deep learning approach has been applied to this area and most of the models conducts the matching through creating a hierarchical matching structure built on convolutional neural nets (ConvNets). The architecture of our model for measure patient pairs is presented in Figure 1. The models based on ConvNets learn to map input

patient representation to vectors, which can then be used to compute their similarity. These are then used to compute a patient similarity score, which together with the representation vectors are joined in a single representation.

In the following we describe how the intermediate representations produced by the ConvNets model can be used to compute patient similarity scores and give a brief explanation of the remaining layers, e.g. hidden and softmax, used in our network.

Single Convolution feature maps: The aim of the convolutional layer is to extract effective patterns, i.e., discriminative medical concept sequences found within the input record that are common throughout the training instances. In general, let $\mathbf{x}_i \in \mathbb{R}^d$ be the d -dimensional event vector corresponding to the i -th time items. A one-side convolution operation involves a filter $w \in \mathbb{R}^{d \times h}$, which is applied to a window of h event features to produce a new feature. For example, a feature c_i is generated from a window of events $x_{i:i+h-1}$ is defined by:

$$c_i = f(\mathbf{w} \cdot x_{i:i+h-1} + b) \quad (5)$$

where $b \in \mathbb{R}$ is a bias term and f is a non-linear function (we use rectification (ReLU)).

Pooling: The output from the convolutional layer (passed through the activation function) are then passed to the pooling layer, whose goal is to aggregate the information and reduce the representation. This filter is applied to each possible window of features in the event matrix $\{x_{1:h}, x_{2:h+1}, \dots, x_{n-h+1:n}\}$ to produce a feature map $\mathbf{c} = [c_1, c_2, \dots, c_{n-h+1}]$, where $\mathbf{c} \in \mathbb{R}^{n-h+1}$. We then apply a max pooling over the feature map and take the average value $\hat{c} = \max\{\mathbf{c}\}$. The idea is to capture the most important feature one with the highest value for each feature map.

Matching Matrix: Given the output of our basic for processing patient records, their resulting vector representations \mathbf{x}_a and \mathbf{x}_b , can be used to compute a record-record similarity score. We follow the approach of [20] that defines the similarity between \mathbf{x}_a and \mathbf{x}_b vectors as follows:

$$\text{sim}(\mathbf{x}_a, \mathbf{x}_b) = \mathbf{x}_a^T \mathbf{M} \mathbf{x}_b^T \quad (6)$$

where $\mathbf{M} \in \mathbb{R}^{m \times m}$ is a similarity matrix. The similarity matrix \mathbf{M} is a symmetrical parameter of the network and is optimized during the training.

Softmax: The output of the penultimate convolutional and pooling layers is flattened to a dense vector \mathbf{x} , which is passed to a fully connected softmax layer. It computes the probability distribution over the labels.

E. Optimization

For different tasks, we need to utilize different loss functions to train our model. Taking regression as an example, we can use square loss for optimization:

$$\mathcal{L}(S_1, S_2, y) = (y - M(S_1, S_2))^2 \quad (7)$$

where $y \in R$ is the real-valued ground-truth label to indicate the matching degree between S_1 and S_2 .

All parameters of the model, including the parameters of word embedding, neural tensor network, spatial RNN are jointly trained by BackPropagation and Stochastic Gradient Descent. Specifically, we use AdaGrad [21] on all parameters in the training process.

Regularization For regularization we employ dropout on the penultimate layer. Dropout prevents co-adaptation of hidden units by randomly dropping out i.e., setting to zero a proportion p of the hidden units during forwardbackpropagation.

IV. EXPERIMENTS AND EVALUATION

In this section, we evaluate our framework on a real clinical EHRs dataset. We carry out the cohort studies by selecting several chronic diseases associated with a range of comorbidities. There are some reasons for our cohort selection. First, they are frequently occurred diseases being extensively analyzed in healthcare applications. Second, these diseases are highly associated with each other, and their combination presents many diagnostic challenges. More importantly, due to the long period progression path of those disease, there are a great deal of temporal information embedded in the medical events. Many of medical research based on machine learning ignored the temporality while our model effectively extract those features and enrich the patients representations. Based on patients clinical similarities derived from these representations, we group patients into clusters by some classical clustering algorithms. As we focus on matching similar patients, the clustering evaluations verify the effectiveness of our model.

As testing our model on the real world EHRs, we demonstrate that our method can effectively represent patients without sacrificing temporal information. With the distributional continuous representations, we apply deep neural networks to derive measure of similarities amongst patients in the datasets.

We then make use of the similarity matrix to group patients. For the evaluations shown in the results, we are convinced that the deep medical event embedding achieves a significant improvement in patients representations.

Further more, we demonstrate the robustness of our model in the cohort studies. As mentioned in [22], the primary disadvantage of medical cohort study is the limited control the investigator has over data collection. The existing data may be incomplete, inaccurate, or inconsistently measured between subjects [23]. As a result, we process patients EHRs for constructing two kinds of data sets. One covers the whole complete patient events for global features analyzing. On another data set, we remove particular events labeled as cohort identifiers from patients EHRs to provide more natural setting in clinical cases. We systematically analyze the performance of our model in the above two settings, and draw some conclusions through our result discussions.

A. Datasets

Our model is trained on a real world longitudinal EHRs database of 218,680 patients for the course of over four years. According to the reasons presented at the beginning of this section, we select four patient cohorts from the EHRs data, namely, Chronic Obstructive Pulmonary Disease (COPD), Diabetes, Heart Failure, and Obesity.

Table I provides a summary of the patient cohorts used in our experiment. Each cohort consists of a set of case patients who are confirmed with one of the four diseases according to their medical diagnosis, and each patient comes with a set of medical events including diagnosis and medications. In each patient encounter, we use the International Classification of Disease-Version 9 (ICD-9) codes to denote the diagnosis of diseases that a patient suffers from. All the clinical events about medications are pre-processed to normalize the descriptions based on brand names and clinical dosages.

Cohorts	# Patients	# Events
COPD	2,000	247,043
Diabetes	2,000	259,074
Obesity	2,000	211,496
Heart Failure	1,135	165,254
Total	7,135	882,867

Table I: Summary of EHRs datasets for patients clustering.

We construct datasets with medical events collected from patients who were confirmed of having the disease by medical experts. We develop the criteria that any patients presented in the datasets has at least forty events. The requirement is set to ensure that each test case has minimum events of clinical history that could be used in reasonable analytics tasks in healthcare. Also, to enable distinctly cluster without overlapping among cohorts, we remove patients who suffers from more than one disease in the cohort list. Finally, there are 8,000 remaining patients and 6,064 distinct clinical events. Medical event appearing in more than 90% of patients or

present in fewer than five patients are removed from the datasets to avoid biases and noise in the learning process.

In the following experiments, we use two datasets: DATASET-I uses the complete patients events while DATASET-II reserves historical events except those labeled as cohort identifiers. On DATASET-I, we split the dataset into training and test sets with same number of patients, and other patients left for validation. As for DATASET-II, we construct the data sets in accordance with DATASET-I. A few of patients are filtered out because of the limited number of their medical events. Table II summaries the two datasets.

Data	# Patients	# Events
DATASET-I		
TRAIN	3,211	396,072
TEST	3,210	399,804
DEV	714	86,991
DATASET-II		
TRAIN	3,083	373,145
TEST	3,080	377,287
DEV	685	81,392

Table II: Summary of modeling datasets.

B. Medical concept embeddings

We use word embeddings to represent each medical event as a vector. We run `word2vec` on the datasets containing 218,680 patients with around 16.9 million medical event records. To learn the embeddings, we choose the Bag of Words model with window size setting to 20 and events filtering with frequency less than 5. The dimensionality of our embedding vectors d is set to 20, 30, 50, 200, 500, respectively, for the comparison purpose, and after a serial practices we select 50 as medical event dimension according to the best performance. Finally, the resulting event matrix covers around 8,000 events which are presented using 50-dimensional vectors, and the event matrix contains all of medical features of patient. Next, we shall discuss how to use them for representing individuals and measuring their distances.

C. Experimental Settings

The parameters of our deep learning were as follow: the width of the convolution filters w is set to 5, 10, 15, 20, 25, and the number of convolutional feature maps m takes on 50, 100, 150, 200. We use stochastic gradient descent to optimize the model’s parameters. We train the model with 50 examples of shuffled mini-batches. We adopt non-linear rectification (ReLU) activation function and a simple max-pooling to achieve the intermediate representations. With regards to overfitting issue we add dropout regularization with dropout rate setting to 0.5.

To optimize our deep features embedding, we conduct experiments using several different parameters sets $\theta = \{d, w, m\}$, which vary in size of `word2vec` embedding dimension, convolution filters width, and the number of convolutional feature maps. In order to find optimal set of parameters, we compare the performance of clustering with only one variable of d, w, m varies.

We implement the clustering base on following representations: (1) One-hot representation. Patient is represented as an event matrix. The matrices are composed of medical event columns, the dimension of which is set to 8,000, or the number of distinct medical events. The event matrix is naturally sparse, but it simplifies patients descriptions. (2) “Shallow” embeddings. As described in section IV-B, we make progress in patients representations with medical event embedding by `word2vec`. Similar to One-hot representation, we represent patients as matrices, but denser and lower dimensional. The dimension of matrix columns has been reduced, with setting from 50 to 800. (3) Deep embeddings. To achieve a deep representation, we combine CNN with distributional medical events embeddings from `word2vec`. Based on above event matrix representations, patients features are filtered through the convolutional layer of neural network. Feature maps that represent patients clinical characteristics are then used to measure patients distances.

With generated representations of each patient, we firstly calculate the similarity amongst all the test patients. Then, we group patients cohorts by matching pairs of patients according to their similarity. Kmeans and Active PCKMeans [24] are adopted for grouping patients based on the first two representations. Also, we compare our model with another metric learning algorithm that have shown state-of-the-art results in clustering. Besides deep neural networks we have applied to learn patients features, we present other two unsupervised methods for calculating patients distances as complementary. Specifically, we use RV and $dCov$ coefficient to calculate correlations of patient feature matrices what derived from `word2vec` embedding.

We verify the cohort discover studies by evaluating the clustering using three popular criteria: *Rand index*, *Purity* and *normalized mutual information(NMI)*.

Rand index is frequently used in data clustering, it is computed as following in [25]:

$$RI = (TP + TN) / \binom{n}{2}$$

where TP counts the number of right decision we have made on grouping pairs of patients who are in the same cohort into one cluster, TN is the number of pairs of patients who came from different cohorts are grouped into dissimilar categories. In general, bad clusterings have RI values close to 0, a perfect clustering has a RI of 1.

Purity is one of primary validation measure to evaluate the cluster quality. We compute *Purity* as defined in [26]:

$$Purity(Cluster, Cohort) = \frac{1}{N} \sum_I \max_j |p_i \cap q_j|$$

where $Cluster = \{p_1, p_2, \dots, p_I\}$ is the set of clusters, $Cohort = \{q_1, q_2, \dots, q_J\}$ is the group of classes, or cohorts in our case. The cohort is identified by the categories of dominant patients in cluster. Similar to *Rand index*, the *Purity* has upper bound of 1 corresponding to the perfect match between the partitions and lower bound of 0 that indicates the opposite.

NMI measures the information shared by the two clusterings, thus can be adopted as a clusterings similarity measure. We follow the form defined in [27] to calculate *NMI* value.

$$NMI(C_{cluster}, C_{ohort}) = \frac{I(C_{cluster}, C_{ohort})}{[H(C_{cluster}) + H(C_{ohort})] / 2}$$

where,

$$I(X, Y) = \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)}$$

is the mutual information between the random variables X and Y ,

$$H(X) = - \sum_{x \in X} p(x) \log p(x)$$

is the information entropy of a discrete random variable X . $p(x)$, $p(x, y)$ are the probabilities of a object being in cluster X and in the the intersection of X and Y . The *NMI* has a fixed lower bound of 0 and upper bound of 1. In our case, $NMI(C_{cluster}, C_{ohort})$ takes its maximum value of 1 when grouping clusterings are identical to the real cohorts, if the partition found is totally independent of the real cohorts, then $NMI(C_{cluster}, C_{ohort}) = 0$.

There are other popular measures for cluster evaluation, namely, *Precision* [28], *Recall*, and *F-measure* [29]. We also validate results by these evaluation.

D. Results and discussion

1) *Performance Comparison*: Table III summaries the results of clustering. As we can see, the deep model with feature embedding is clearly superior to others. On DATASET-I, the deep embedding model achieves an average *Rand index* of 0.9887, comparing with the second best one with 0.6796. Measured by *Purity* and *NMI*, it can achieve the performances of 0.9882 and 0.9516, separately, which also outperforms others with a margin. The superiority of the model is illustrated in DATASET-II as well, which is a more difficult task. Measured by *Purity* and *NMI*, KMeans¹ and Active PCKMeans¹ achieve 0.3367, 0.0351 and 0.4410, 0.0682 separately. KMeans² and Active PCKMeans² can only improve 11% and 25% on *Purity* respectively. However, our CNN model achieves more than 50% improvement over them.

As a reasonable explanation, we hold that the deep features learning can be viewed as a two-stage model. During the first stage, the clinical features of each patients are summarized in the shallow *word2vec* embeddings, making progress with nearly 10% improvement. Next, global features are learned base on local context features came from *word2vec*. The deep learning representation makes continuous improvement, which leads to a ultimate expression of patients. Figure 2 shows how expressive representations of patients contribute to match patients cohorts. With a significant 48% improvement produced, we illustrate the effectiveness of our deep embedding model in expressively representing patients.

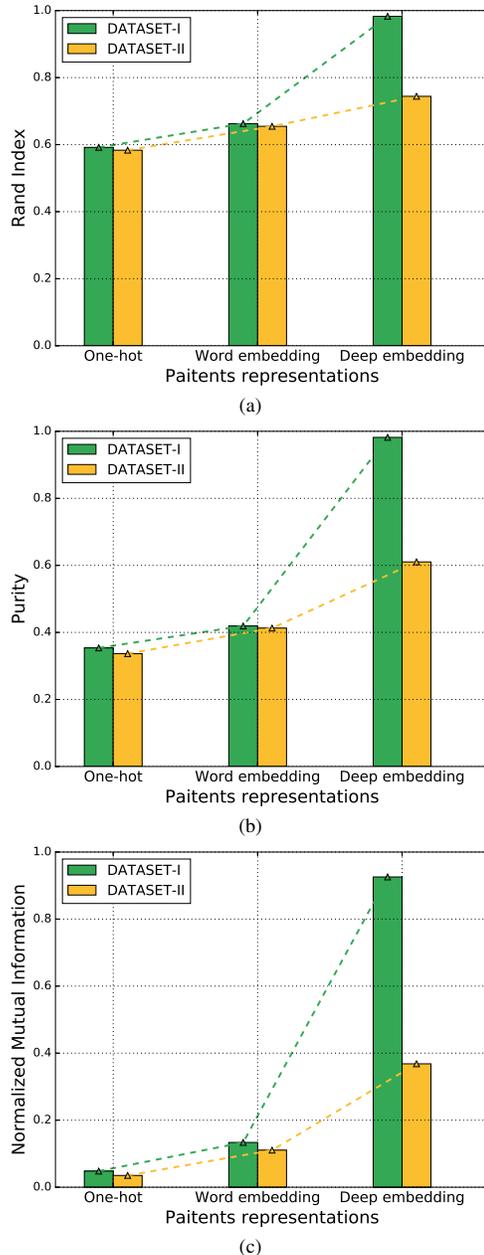


Figure 2: Performance of different representations.

2) *Parameters Optimization*: Figure 3 illustrates the optimizations of hyperparameters in our model. The line charts in one row assess what effects the variation has on grouping patients. As results summarized in Figure 3a, Figure 3b, Figure 3c, the dimension of medical event embedding have little effect on DATASET-I. That because our deep learning model have successfully obtained the primary features in the patients representations, achieving nearly perfect 1 of *RI*, *Purity*, and *NMI*. We make determinations based on DATASET-II. According to the performance lines shown in the figures, three clustering evaluations we choose—*RI*, *Purity* and *NMI* achieves the best performance at the same time, with 50 dimensionality embedding, 100 feature maps, and

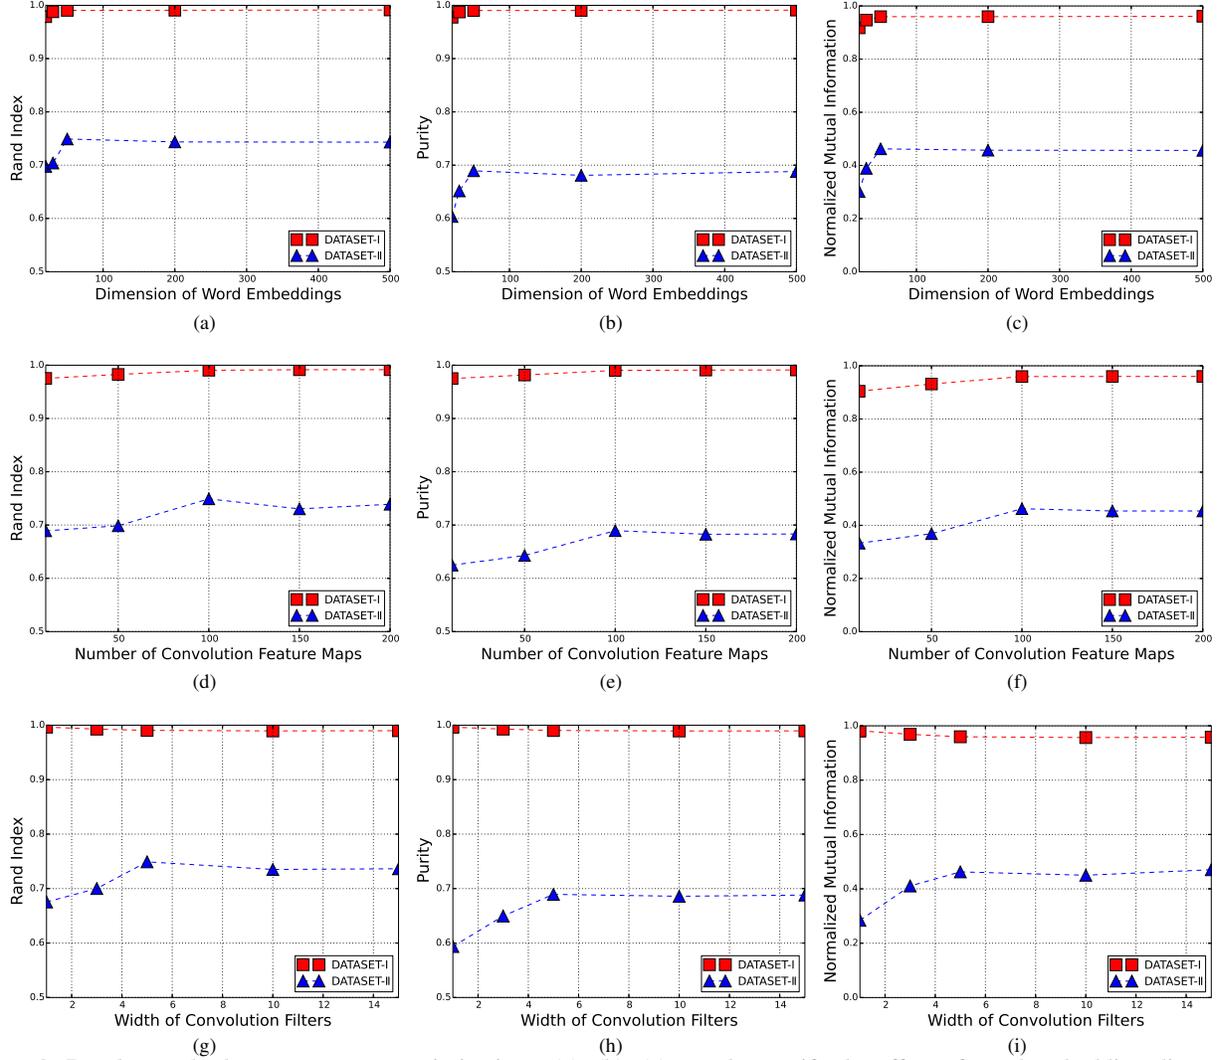


Figure 3: Results on the hyperparameters optimizations. (a), (b), (c) together vary the effect of word embedding dimension on clustering performance. (d), (e), (f) measures the efficacy of variation on convolution feature maps. (g), (h), (i) make determinations of convolution feature filters width.

Method	DATASET-I			DATASET-II		
	<i>RI</i>	<i>Purity</i>	<i>NMI</i>	<i>RI</i>	<i>Purity</i>	<i>NMI</i>
KMeans ¹	0.5919	0.3538	0.0481	0.5834	0.3367	0.0351
Active PCKMeans ¹	0.6506	0.4801	0.0976	0.6451	0.4410	0.0682
KMeans ²	0.6627	0.4192	0.1336	0.6547	0.4129	0.1106
Active PCKMeans ²	0.6796	0.5610	0.1682	0.6794	0.5487	0.1622
Matric Learning	0.6732	0.5103	0.1049	0.6475	0.4013	0.1055
Our model(<i>RV</i>)	0.6679	0.4457	0.2301	0.6285	0.3858	0.1037
Our model(<i>dCov</i>)	0.6708	0.4475	0.2361	0.6268	0.3858	0.1048
Our model(CNN)	0.9887	0.9882	0.9516	0.7491	0.6894	0.4624

Table III: Evaluations of cohort discovering on DATASET-I, DATASET-II. KMeans¹, Active PCKMeans¹ groups patients with One-hot representations, where KMeans², Active PCKMeans² adopt the Shallow embeddings to match similar patients pairs. We contrast the performance of our model at DATASET-I, DATASET-II with the same parameters. Values of *RI*, *Purity*, *NMI* presented in the table are average of a group of results.

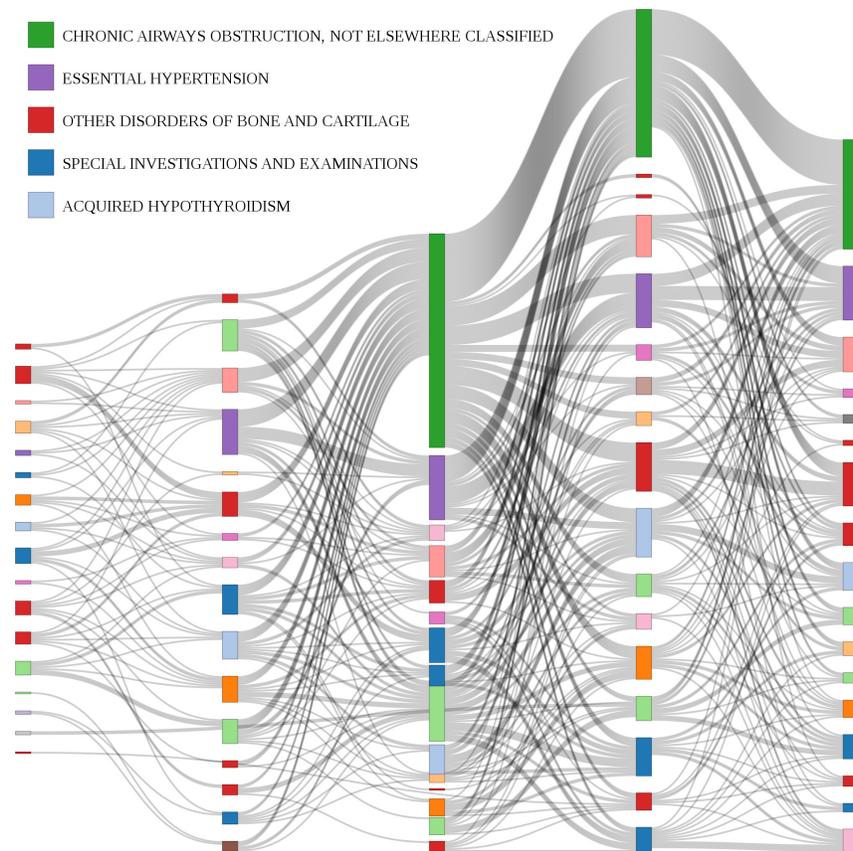


Figure 4: An Example of sankey-pathway on the COPD dataset. Each color represents different features and only top frequent feature names are listed.

5 convolution filter width. The consistent performances of different measures assessed in our experiments show that the optimizations of parameters are correct and free of bias.

3) *Discussions:* Table III provides the comparisons of clustering results on DATASET-I and DATASET-II. The results of deep embedding on DATASET-II exhibit a steady outperformance over other methods. On DATASET-II, the deep embedding model is trained with fewer medical events than DATASET-I. As expected, the evaluations of identifying patient cohorts is slight affected by the data we deal with. Compared to the sterling performances on DATASET-I, the *RI* resulted in DATASET-II drops to 0.75 with loss of 0.23. One simple but reasonable explanation is that the events removing from dataset cause a loss of many a medical features. Even though, our model extract remain features effectively offering promising performance. To sum up, we verify that our deep learning method works effectively in representing patients and learning specific features that are not present or missing.

Table III also reports the comparisons on effectiveness of our supervised and unsupervised measurement of patient similarity. On average, the unsupervise measurements—*RV* and *dCov*, respectively gets *RI* of 0.67 and *NMI* of 0.23, which are 31% and 72% lower than the deep learning model(0.97). Although, it's worthwhile to mention that our unsupervise

models achieve at least 12% improvement over the baseline. Comparing to the semi-supervise method *Active PCKMeans* proposed in [24], our model achieves the same performance but do not need training examples. These comparisons suggest that our models consistently and significantly surpass other patient representations.

4) *Visual Analysis:* We make a study on medical events sequence mining for representation. To discovery the medical pattern hidden behind the EHRs about COPD, we select top-100 similar patients from the COPD cohort , whom are grouped into true cluster by our method. We extract the common events occurred in the EMR of many patients. The SanKey diagram presents the progression path of the medical events collected from patients EHRs. As shown in the Figure 4, the green, purple and red bar are related closely, which respectively represent Chronic Airways Obstruction, Essential Hypertension, Other Disorder of Bone and Cartilage. The interactions of those diseases presented in the diagram has been validated by lots of medical research in the real world, that convinces us of the applicability of our model.

In summary, the results of experiments clearly demonstrate the effectiveness of the deep model with medical feature embedding on real EHRs data. Theoretically, our model benefits from the large number of convolutional filters and lower event

embedding dimensionality. It is notable that our model has several important hyper-parameters like word2vec window size, dimensionality of clinical event vector, the number of convolutional filters. To be more realistic, we narrow down the scopes of variations and select the best performance values.

V. CONCLUSIONS

Patient similarity assessment is the enabling technique for various healthcare applications, such as disease subtyping and evidence based medicine. However, due to the complexity of medical data, extracting effective patient representations confronts distinct challenges. Though useful, most existing models proposed to discover hidden patterns in EHRs overlook the temporal information of medical events. In this paper, we propose a deep learning framework to learn patient representations for similarity measuring, in which the temporal properties of EHRs are preserved. The experimental results show that our model achieves significantly better representations over the baselines, which enables more accurate patient cohort discovery. Our next plans include solving the data irregularity issue by adding the time interval information and applying this techniques in other domain, such as health visualization. Besides, it can be observed in the experiments that our unsupervised scheme also succeeded in matching similar patients.

VI. ACKNOWLEDGEMENT

This work is sponsored by "The Fundamental Theory and Applications of Big Data with Knowledge Engineering" under the National Key Research and Development Program of China with grant number 2016YFB1000903; National Science Foundation of China under Grant Nos.61428206; Ministry of Education Innovation Research Team No.IRT13035. The paper is also supported by Project of China Knowledge Centre for Engineering Science and Technology. The work of F. Wang is partially supported by National Science Foundation under Grant Number IIS-1650723

REFERENCES

- [1] Y. Cheng, F. Wang, P. Zhang, and J. Hu, "Risk prediction with electronic health records: A deep learning approach," 2016.
- [2] J. Sun, F. Wang, J. Hu, and S. Ebadollahi, "Supervised patient similarity measure of heterogeneous patient records," *ACM SIGKDD Explorations Newsletter*, vol. 14, no. 1, pp. 16–24, 2012.
- [3] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [4] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [5] J. Josse and S. Holmes, "Measures of dependence between random vectors and tests of independence. literature review," *arXiv preprint arXiv:1307.7383*, 2013.
- [6] G. J. Székely, M. L. Rizzo, N. K. Bakirov *et al.*, "Measuring and testing dependence by correlation of distances," *The Annals of Statistics*, vol. 35, no. 6, pp. 2769–2794, 2007.
- [7] L. Chan, T. Chan, L. Cheng, and W. Mak, "Machine learning of patient similarity: A case study on predicting survival in cancer patient after locoregional chemotherapy," in *Bioinformatics and Biomedicine Workshops (BIBMW), 2010 IEEE International Conference on*. IEEE, 2010, pp. 467–470.
- [8] F. Wang, J. Hu, and J. Sun, "Medical prognosis based on patient similarity and expert feedback," in *Pattern Recognition (ICPR), 2012 21st International Conference on*. IEEE, 2012, pp. 1799–1802.
- [9] F. Wang, N. Lee, J. Hu, J. Sun, and S. Ebadollahi, "Towards heterogeneous temporal clinical event pattern discovery: a convolutional approach," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012, pp. 453–461.
- [10] K. Ng, J. Sun, J. Hu, and F. Wang, "Personalized predictive modeling and risk factor identification using patient similarity," *AMIA Summits on Translational Science Proceedings*, vol. 2015, p. 132, 2015.
- [11] N. Kasabov and Y. Hu, "Integrated optimisation method for personalised modelling and case studies for medical decision support," *International Journal of Functional Informatics and Personalised Medicine*, vol. 3, no. 3, pp. 236–256, 2010.
- [12] M. J. Sewitch, K. Leffondré, and P. L. Dobkin, "Clustering patients according to health perceptions: relationships to psychosocial characteristics and medication nonadherence," *Journal of psychosomatic research*, vol. 56, no. 3, pp. 323–332, 2004.
- [13] R. Henao, J. Murray, G. Ginsburg, L. Carin, and J. E. Lucas, "Patient clustering with uncoded text in electronic medical records," in *AMIA Annual Symposium Proceedings*, vol. 2013. American Medical Informatics Association, 2013, p. 592.
- [14] G. T. Huang, K. I. Cunningham, P. V. Benos, and C. S. CHEN-NUBHOTLA, "Spectral clustering strategies for heterogeneous disease expression data," in *Pacific Symposium on Biocomputing, Pacific Symposium on Biocomputing*. NIH Public Access, 2013, p. 212.
- [15] D. Kiela and L. Bottou, "Learning image embeddings using convolutional neural networks for improved multi-modal semantics," in *EMNLP*. Citeseer, 2014, pp. 36–45.
- [16] A. Severyn and A. Moschitti, "Learning to rank short text pairs with convolutional deep neural networks," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2015, pp. 373–382.
- [17] F. Kherif, J.-B. Poline, S. Mériaux, H. Benali, G. Flandin, and M. Brett, "Group analysis in functional neuroimaging: selecting subjects using similarity measures," *NeuroImage*, vol. 20, no. 4, pp. 2197–2208, 2003.
- [18] B. Hu, Z. Lu, H. Li, and Q. Chen, "Convolutional neural network architectures for matching natural language sentences," in *Advances in Neural Information Processing Systems*, 2014, pp. 2042–2050.
- [19] Z. Lu and H. Li, "A deep architecture for matching short texts," in *Advances in Neural Information Processing Systems*, 2013, pp. 1367–1375.
- [20] A. Bordes, J. Weston, and N. Usunier, "Open question answering with weakly supervised embedding models," in *Machine Learning and Knowledge Discovery in Databases*. Springer, 2014, pp. 165–180.
- [21] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *The Journal of Machine Learning Research*, vol. 12, pp. 2121–2159, 2011.
- [22] J. W. Song and K. C. Chung, "Observational studies: cohort and case-control studies," *Plastic and reconstructive surgery*, vol. 126, no. 6, p. 2234, 2010.
- [23] W. S. Browner, S. B. Hulley, and S. R. Cummings, *Designing clinical research: an epidemiologic approach*. Lippincott Williams & Wilkins, 1988.
- [24] S. Basu, A. Banerjee, and R. J. Mooney, "Active semi-supervision for pairwise constrained clustering," in *SDM*, vol. 4. SIAM, 2004, pp. 333–344.
- [25] W. M. Rand, "Objective criteria for the evaluation of clustering methods," *Journal of the American Statistical association*, vol. 66, no. 336, pp. 846–850, 1971.
- [26] C. D. Manning, P. Raghavan, H. Schütze *et al.*, *Introduction to information retrieval*. Cambridge university press Cambridge, 2008, vol. 1, no. 1.
- [27] M. Meilä, "Comparing clusterings an information based distance," *Journal of multivariate analysis*, vol. 98, no. 5, pp. 873–895, 2007.
- [28] Y. Zhao and G. Karypis, "Criterion functions for document clustering: Experiments and analysis," Citeseer, Tech. Rep., 2001.
- [29] C. J. Van Rijsbergen, "Foundation of evaluation," *Journal of Documentation*, vol. 30, no. 4, pp. 365–373, 1974.